Contents

List of Figures

List of Tables

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **2006**	2. REPORT TYPE	3. DATES COVERED **00-00-2006 to 00-00-2006**
4. TITLE AND SUBTITLE **A Hybrid Collision Avoidance Scheme for Ad Hoc Networks**		5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **University of California at Santa Cruz,Department of Computer Engineering,Santa Cruz,CA,95064**		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR’S ACRONYM(S) 11. SPONSOR/MONITOR’S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT **Approved for public release; distribution unlimited**		
13. SUPPLEMENTARY NOTES		
14. ABSTRACT		
15. SUBJECT TERMS		

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**		**19**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

A Hybrid Collision Avoidance Scheme for Ad Hoc Networks *

Yu Wang J.J. Garcia-Luna-Aceves
Department of Computer Engineering
University of California at Santa Cruz
Santa Cruz, CA 95064, U.S.A.
Email:{*ywang,jj*}*@cse.ucsc.edu*

Abstract

A novel hybrid collision avoidance scheme that combines both sender-initiated and receiver-initiated collision-avoidance handshakes is proposed for multi-hop ad hoc networks. The new scheme is compatible with the popular IEEE 802.11 MAC protocol and involves only some additional queue management and book-keeping work. Simulations of both UDP- and TCP-based applications are conducted with the IEEE 802.11 MAC protocol, a measurement-based fair scheme and the new scheme. It is shown that the new scheme can alleviate the fairness problem with almost no degradation in throughput. More importantly, it is shown that without explicit information exchange among nodes, the fairness problem cannot be solved conclusively if reasonable throughput is to be maintained.

1 Introduction

Many contention-based channel access schemes have been proposed for multi-hop ad hoc networks in the literature [1–6]. Collision avoidance is very important in these schemes to combat the adverse effects of hidden terminals [7] and can be largely divided into two categories, sender-initiated and receiver-initiated. In sender-initiated schemes, a node with a packet to send initiates the collision avoidance handshake with an intended receiver. Usually the handshake comprises the exchange of short ready-to-send (RTS) and clear-to-send (CTS) control packets between a sender and a receiver followed by the transmissions of the actual data packet and the optional acknowledgment packet. The RTS and CTS packets carry information about the duration of the handshake and serve as a channel reservation scheme to notify overhearing nodes to defer their access to the shared channel to avoid collisions. On the other hand, in receiver-initiated schemes, a node has to poll its neighbors actively to see if they have packets for itself. The rationale behind receiver-initiated schemes is that, a receiver usually has better knowledge of the contention around itself and collision avoidance is more important at the receiver's side as the receiver needs to receive relatively long data packets successfully which are more vulnerable to interference. It has been shown that, if the polled nodes always have packets for the polling node, receiver-initiated schemes with proper collision avoidance procedures can outperform sender-initiated schemes by reducing the overhead of control packets [6]. Otherwise, the performance may degrade due to wasted transmissions of polling packets that poll inactive nodes with no packets for the polling node. The degradation in performance will be more conspicuous in light to medium traffic load, unless a good traffic predictor is available at the polling node.

Despite the potential benefits of receiver-initiated schemes, they have not received wide acceptance. One reason is that sender-initiated schemes are more straightforward, because a sender has full knowledge of the packets in its queue and it can initiate the collision avoidance handshake only when necessary. On the other hand, for receiver-initiated schemes, a good traffic estimator and an appropriate polling discipline that can be adapted to the dynamic environments of ad hoc networks are mandatory and they have not been investigated sufficiently so far. Another reason is the prevalent acceptance of the IEEE 802.11 MAC protocol in the research community, which uses a sender-initiated collision avoidance scheme. Many performance enhancements have been proposed and they are confined to the sender-initiated framework stipulated by the IEEE 802.11 MAC protocol.

*This work was supported in part by the US Air Force/OSR under Grant No. F49620-00-1-0330.

Despite its popularity, the IEEE 802.11 MAC protocol can suffer severe fairness problems in multi-hop ad hoc networks where location-dependent contention is common. As is already pointed out in the research literature [8–11], some nodes in such networks are at a disadvantage in contending with other nodes due to their locations and may suffer severe degradation in throughput. Additionally, the commonly used binary exponential backoff (BEB) scheme, despite its robustness against repetitive collisions, can aggravate the fairness problem, because the node that succeeds in the last transmission period will gain access to the shared channel again with much higher probability while other nodes suffer starvation. Two basic classes of schemes have been proposed in the recent past to solve the fairness problem in 802.11. In the first class, the goal is to achieve max-min fairness [8, 12, 13] by reducing the ratio between maximum throughput and minimum throughput of flows, either at a node's level or at a flow's level. In the second class, the approach used in fair queueing for wireline networks is adapted to multi-hop ad hoc networks, taking into account the salient characteristics of such networks such as location-dependent contention, distributed coordination and possible spatial reuse [9–11, 14, 15]. Flow contention graphs are used extensively in the schemes in this second class to model the contention among nodes. Figure 1 shows an example of how this is done. Any two flows with adjacent vertices in the flow contention graph should not be scheduled to transmit at the same time. All these schemes usually exhibit some form of tradeoff between throughput and fairness. Nodes that are *leading* in channel access (in terms of throughput) will decrease their channel access activities, while nodes that are *lagging* will increase their channel access activities. In this way, nodes are encouraged to compete fairly but at the cost of increased contention, which may degrade the overall throughput. Despite the differences of backoff algorithms and information exchange among these schemes, the underlying channel access scheme remains largely the basic sender-initiated collision avoidance handshake, which can be less effective than a receiver-initiated scheme when a receiver has better knowledge of the contention around itself than the sender.This motivates us to design an adaptive collision avoidance scheme that makes use of both sender-initiated and receiver-initiated handshakes, because a receiver-initiated handshake is more desirable in some cases and a better tradeoff between throughput and fairness may be achieved. To expedite its introduction, the new hybrid scheme should fit within the IEEE 802.11 framework, even though it combines both sender-initiated and receiver-initiated handshake, and nodes implementing the new scheme should not break an existing 802.11-based network. Furthermore, the new scheme should be simple and not introduce new types of control packets, because they may complicate implementation of the finite state machine of the protocol and degrade the overall network throughput unnecessarily when the basic sender-initiated scheme suffices.

The rest of the paper is organized as follows. In Section 2, the new hybrid scheme is specified, which in fact is a very simple extension to the existing IEEE 802.11 MAC protocol and involves only some additional queue management and book-keeping work. In Section 3, a measurement-based fair scheme [12] is described, which is one of the simple and straightforward fair schemes proposed so far; it does not require explicit information exchange among nodes and serves as a basis for comparison with our new hybrid scheme. In Section 4, simulations with the original IEEE 802.11 MAC protocol, the measurement-based fair scheme and the new hybrid scheme are presented for both UDP- and TCP-based traffic. It is shown that various degrees of the fairness problem exist in the original IEEE 802.11 MAC protocol even for simple network configurations with only two competing flows. Although the new hybrid scheme cannot solve the fairness problem conclusively, it can alleviate the fairness problem in some cases with almost no degradation in throughput. It is also reasoned that more explicit information exchange among nodes is mandatory to solve the fairness problem while maintaining reasonable throughput. Section 5 concludes this paper with directions for future work.

2 The New Hybrid Collision Avoidance Scheme

Bharghavan et al. [2] proposed a request-for-request-to-send (RRTS) packet to alleviate some of the interference problems due to hidden terminals in their seminal paper to address the fairness problem. Talucci and Gerla [5] proposed MACA-BI (Multiple Access with Collision Avoidance - By Invitation) which was the first receiver-initiated MAC protocol. Garcia-Luna-Aceves and Tzamaloukas [6] advanced that work and proposed several collision-free RIMA (receiver-initiated multiple access) protocols. Here collision-free means that, once a node sends a data packet, the data packet can be received by the receiver successfully, given that the channel is ideal without impairment and the only cause of failure to receive a packet is concurrent transmissions from multiple nodes. RIMA protocols achieve this collision-free property by introducing some additional types of short control packets and enforcing various collision-avoidance waiting periods. The receiver-initiated handshake in our proposed hybrid channel access scheme is simpler than that in the RIMA protocols. Firstly, it does not introduce new types control packets. Instead, a CTS packet is used as the polling packet to maintain compatibility with the original IEEE 802.11 MAC protocol. Secondly, it does not include the various collision-avoidance waiting periods enforced in RIMA protocols. Instead, nodes defer access to the shared channel according to the network allocation vector (NAV) included in those overheard packets, which specifies the duration of the ensuing handshake. The reason is that the IEEE 802.11 MAC protocol itself cannot ensure collision-free data packet transmissions. We opt not to introduce additional collision-avoidance procedures and try to maintain compatibility with the existing protocols. Hence, the receiver-initiated collision avoidance handshake just includes a three-way CTS-data-ACK exchange between polling and polled nodes. Though it is not expected that the hybrid scheme will improve throughput because it does not provide strict collision avoidance, it may still alleviate the fairness problem, because both a sender and a receiver can initiate a collision avoidance handshake alternately and the burden of contending for the shared channel is distributed to participating nodes according to the different degrees of contention they experience.

Our hybrid collision avoidance scheme is built around the framework of the IEEE 802.11 MAC protocol. A node that implements this scheme operates alternately in two modes, sender-initiated (SI) and receive-initiated (RI). The SI mode is the default mode, which is in effect the same as the original IEEE 802.11 MAC protocol. The usual four-way RTS-CTS-data-ACK handshake is used in the SI mode. The aforementioned receiver-initiated three-way collision avoidance handshake is used in the RI mode introduced in the hybrid scheme. , and is triggered only when the SI mode does not perform well. In this mode, more cooperation between a pair of sending and receiving nodes is required, because both of them need to enter the RI mode before the receiver-initiated handshake can be initiated.

The only necessary change to the frame structures in the IEEE 802.11 standard to implement the hybrid scheme is the addition of the RI flag. Figure 2 illustrates the frame structure of the IEEE 802.11 RTS frame (ref. Fig. 13 in Page 35 and Fig. 16 in Page 41 of the IEEE 802.11 standard [3]). Given that the *More data* bit is not used in the ad hoc mode according to the standard, it may be reused as the RI flag to indicate if the RI mode is on or not. Nodes that do not implement the hybrid collision avoidance scheme can safely ignore this bit. The states of both sending and receiving nodes in the proposed hybrid collision avoidance scheme are shown in Figure 3 and are explained separately. A sender enters *RI setup* mode when it sends the same RTS packet for more than one half of the times allowed in the IEEE 802.11 MAC protocol and has no response from the intended receiver. Failure to obtain a response from the intended receiver usually implies that contention around the receiver is so severe that the receiver is prevented from responding. Hence, it is more appropriate to let the receiver start the collision-avoidance handshake when this happens. The number of unsuccessful RTS packets to activate the RI mode is chosen according to the simulations with the network configurations investigated in this paper. The present value in the scheme is able to achieve balanced results while other values may perform very differently for different configurations.

After the sender enters *RI setup* mode, it sets the RI flag in all the subsequent RTS packets and other packets that it sends out and requests the intended receiver to enter the RI mode as well. During this stage, the node keeps sending RTS packets following the usual collision-avoidance procedures, because it has not established an *association* with the intended receiver. There are two possible outcomes. One outcome is that the node never gets any CTS packet from the intended receiver. In this case, the sender may declare the receiver down after it has to drop a few packets. The other outcome is that it receives a CTS packet from the intended receiver. In this case, the sender enters the *RI-associated* mode and will not send an RTS to the receiver thereafter. This helps to reduce the contention around the receiver and also makes the sender available for accepting polling requests from the receiver. To keep the receiver in the RI mode, the sender keeps setting the RI flag in all the data packets that it sends out. The RI flag is cleared only when the sender's queue becomes empty.

The receiver enters and stays in the RI mode when it receives RTS packets or data packets destined to it with the RI flag set. The receiver then generates RI-response packets (which are in fact self-initiated CTS packets) and multiplexes them with other data packets in its MAC queue. However, the receiver should not generate RI-response packets indiscriminately when it receives a packet with the RI flag on, lest serious fairness problem may occur. This can be explained as follows. When an RI-response packet becomes the head-of-line (HOL) packet of a receiver's queue, the node will send a self-initiated CTS to the sender, which in fact serves as the ready-to-receive (RTR) packet to poll the sender in the RIMA protocols [6]. If the sender replies with a data packet with the RI flag still on, which implies that there are more packets in its sending queue, the receiver will add another RI-response packet to the end of its queue. If there is no packet for other nodes intervened in the MAC queue, the receiver will be *locked into* the sender and will keep sending CTS packets to it. In this way, they may monopolize the shared channel for a long time, which obviously defeats the purpose of the hybrid scheme. Hence, when a node receives a packet with the RI flag on, it checks its HOL packet to see whether it is an RI-response packet for the node that just sent this packet. If so, the RI request is ignored; otherwise, it is added to the end of its MAC queue.

The RI-response packets are treated like RTS packets for normal data packets. That is, when they are served via a successful receiver-initiated CTS-data-ACK handshake or when they are transmitted more than the times allowed for RTS packets in the IEEE 802.11 standard, they will be removed from the MAC queue. Such precautions are necessary. One reason is to avoid excessive delay or deadlock when the sending node is down or moves out of range of the receiving nodes. Another reason is to promote fairness so that neighboring nodes may still get chances to initiate handshake with the receiver or other nodes.

The above specification clearly shows that, with some additional queue management and book-keeping work, the existing IEEE 802.11 can be easily extended to support a receiver-initiated scheme while maintaining compatibility.

3 Measurement-based Fair Scheme

In this section, we describe the measurement-based fair scheme [12] with which we compare the IEEE 802.11 MAC protocol and the new hybrid scheme. The rationale behind the scheme is surprisingly simple. Whenever a node sends or receives a packet, it updates its own estimation of its share (W_{ei}) or other nodes' share (W_{eo}) of the channel depending on the purpose of the packet. To avoid any explicit information exchange among these nodes, each node just treats all the nodes around itself as a single entity that competes against itself. For example, if a node sends an RTS packet, it will update W_{ei} because the RTS packet serves to reserve the channel for itself. If the node receives an RTS packet addressed to itself, it updates W_{eo} because the RTS packet serves for other nodes. Details on the updating of W_{ei} and W_{eo} can be found in [12] and are not repeated here. Then the ratio between W_{ei} and W_{eo}, which is denoted by FI_e, serves as a fairness index to show whether a node is leading or lagging in channel access. If FI_e for a node is larger than a pre-defined constant C ($C > 1$), which implies that the node has obtained more than its fair share, the node doubles its contention window (CW) from which the backoff timer is derived. If FI_e lies between $1/C$ and C, then the node just holds on to its current CW as it estimates that both its neighbors and itself have obtained roughly equal share. If FI_e is smaller than $1/C$, then the node cuts its CW to a half to contend more vigorously for the channel. It should be noted that CW is bounded by the minimum and maximum values stipulated in the IEEE 802.11 standard. The measurement-based fair scheme is shown to be quite effective in the configurations investigated in [12] by sacrificing some throughput for better fairness.

However, this scheme may encounter the problem of severe throughput degradation in some cases, e.g., when two neighboring nodes are engaged in TCP-based connections. This can be explained as follows. In the measurement-based scheme, a node at one end of a TCP connection continuously estimates its share and other node's share of the channel including the node on the other side of the connection. When this node sends one or a few data packets, it estimates that its use of the channel has exceeded its fair share and will increase its contention window accordingly. The node at the other end of the TCP connection behaves similarly. In this way, both nodes may have a contention window that is larger than necessary and the throughput is degraded due to the increased time wasted in waiting. The degradation in throughput can also happen in UDP-based traffic in which two nodes take turns in channel access according to the their own measurements. However, this phenomenon can be more conspicuous in TCP-based connections, because flow control in TCP may also be activated, which can further slow down the channel access activities unnecessarily.

4 Simulation Results

In our simulations, we focus on how two competing flows share the available channel resource in a few simple network configurations. These configurations are shown in Figures 4 and 5, in which a dashed line means that two nodes can hear each other's transmissions and an arrow indicates an active flow between two nodes. Nodes without any line in-between are hidden from each other. These network configurations have the same *flow contention graph* that is commonly used in schemes that try to approximate wireline fair queuing in ad hoc networks, e.g., [11, 13–15]. The following results indicate that various degrees of fairness problem can occur even for such simple networks with the same flow contention graph.

We use GloMoSim 2.0 [16] as the network simulator and implement both the measurement-based fair scheme and the new hybrid scheme based on its implementation of the IEEE 802.11 MAC protocol for fair comparison. Direct sequence spread spectrum (DSSS) parameters are used throughout the simulations, which are shown in Table 1. The raw channel bit rate is 2Mbps.

We investigate the performance of the IEEE 802.11 MAC protocol, the measurement-based fair scheme (for simplicity, it is called MFS thereafter) and the new hybrid scheme under both UDP- and TCP-based traffic. In the first set of simulation experiments, there are two competing UDP-based flows. For each flow, one node keeps sending data packets to the other at a constant bit rate, such that the sending queue is always non-empty. UDP is the underlying transport layer, thus no acknowledgment packet is sent back to the initiating node. Simulation results are shown in Tables 2 and 3.

Table 2 shows the performance of the original IEEE 802.11 MAC protocol. For configurations 4-1 and 4-8, some nodes are almost denied access to the shared channel and suffer severe degradation in throughput. For other configurations, the original MAC protocol works fairly well. Table 3 shows the performance of the original IEEE 802.11 MAC protocol, the MFS scheme and the hybrid scheme. It is apparent that, the hybrid scheme is the same as as the original IEEE 802.11 when the RI mode is not triggered in some network configurations. For simplicity, the performance of all schemes in these configurations is not shown here, because the conclusions to be drawn are unaffected. From Table 3, it is clear that the fairness problems in configurations 4-1 and 4-8 are alleviated significantly, without a sacrifice in throughput when the hybrid scheme is used. In addition, even when the RI mode is triggered unnecessarily in the three other configurations, it has almost no negative effect on throughput. On the other hand, the MFS works best in configurations 4-1 and 4-8. However, in other configurations in which the original IEEE 802.11 MAC protocol works well, the use of MFS degrades throughput unnecessarily, in some cases to less than 50% of the original. It is also worth noting that the aggregate throughput in all these network configurations remains almost the same despite the fact that the fairness problem exists in some configurations. This shows the importance of considering the underlying topologies of networks even though they have the same flow contention graph.

In the second set of simulation experiments, there are two competing TCP-based flows. We use the FTP/Generic application provided in GloMoSim, in which a client simply sends data packets to a server without the server sending any control information back to the client other than the acknowledgment packets required by TCP. Whenever a packet indicates success of delivery by the transport layer (TCP), the client sends the next data packet. It should be noted that the acknowledgment packet from TCP is still regarded as a normal data packet from the view of MAC layer. Hence, due to the peculiarities of the application, it is disadvantageous for the MAC layer to transmit more than one packet at a time. When this is applied to the hybrid scheme, it means that it is more desirable for a node and its peer to leave RI associated mode just after a CTS-Data-ACK handshake is done, so that they can switch the roles of sender and receiver timely. So in the implementation of the hybrid scheme, we make the necessary changes to take this into account. That is, for such type of traffic, a node clears its RI flag when it receives acknowledgment packet from the node that it is sending data packets to. Simulation results are shown in Tables 4 and 5.

It is clear from Table 4 that the fairness problem is much more severe for two competing TCP-based flows than for the case of UDP-based flows. In some cases, such as configurations 4-1, 4-7 and 4-8, one FTP flow is denied access to the shared channel for most of the time. Here throughput of zero does not mean that TCP connection is not set up. Instead, it is because of the extremely low throughput (on the order of a few kilobytes per second) for these flows that the statistics are not shown here. When the hybrid collision avoidance scheme is used, in some cases it is triggered and performs almost the same as the original IEEE 802.11 MAC protocol while in some other cases it is not triggered at all. For simplicity Table 5 shows only the results when there exist differences between these two schemes. It is clear that the hybrid scheme performs slightly better than the original 802.11 MAC scheme for configurations 3-3 and 4-2, while it performs much better in terms of fairness in configuration 4-3, though there is about an 8% degradation in throughput. It is more difficult to improve fairness of TCP-based flows than UDP-based flows due to the flow control and congestion avoidance functions of TCP. A node that suffers excessive packet loss or delay decreases its sending rate according to TCP, which can aggravate the fairness problem already existing at the MAC layer. In such cases, even the hybrid scheme can lose its effectiveness.

On the other hand, the MFS achieves very good fairness in all these configurations but at the cost of much reduced throughput except for configuration 4-7. This is due to the fact that nodes are slowed down to encourage fair contention for the shared channel.

From the radically different results of these networks that share the same flow contention graph, it can be reasoned that the proposed fair schemes that are based on flow contention graphs (e.g., [11, 14, 15]) are not sufficient to solve the fairness problem conclusively. Simple schemes like the MFS cannot solve the fairness problem desirably as well if reasonable throughput is to be maintained. Additionally, though the new hybrid scheme has better throughput and fairness properties, it is also inadequate to solve the fairness problem because it does not make any modification to the existing (unfair) backoff algorithm in the IEEE 802.11 MAC protocol and does not rely on additional contention information, which may be useful for nodes to contend for the shared channel more efficiently and fairly. Hence, more explicit information exchange among nodes as well as the good use of such information should be studied to address both fairness and throughput adequately.

5 Conclusion

In this paper, we have proposed a new hybrid channel access scheme that includes sender-initiated and receiver-initiated collision avoidance. This is based on the observation that sometimes a receiver-initiated scheme is more appropriate when receivers are more knowledgeable of the contention around themselves and can compete for the channel more effectively. By adaptively sharing the burden of initiating the collision-avoidance handshake between the nodes that experience different levels of contention, better fairness may be achieved with almost no degradation in throughput. An attractive feature of the new scheme is that it is a simple extension to the existing IEEE 802.11 MAC protocol and maintains compatibility with the standard. Simulations are conducted with the original IEEE 802.11 MAC protocol, a measurement-based fair scheme (MFS), and the new scheme. It is shown that, although the proposed hybrid scheme does not solve the fairness problem conclusively, it does alleviate the fairness problem in some cases without sacrificing much throughput and simplicity. Simple schemes such as the MFS can achieve far better fairness but sacrifice too much throughput. Hence, it is reasoned that more explicit information exchange among nodes is needed to solve the fairness problem conclusively. Integration of the new hybrid scheme with recently proposed mechanisms [11, 14, 15] that try to approximate fair queueing for ad hoc networks and the use of more contention information in channel access should be addressed in future work.

References

[1] P. Karn, "MACA - a New Channel Access Method for Packet Radio," in *ARRL/CRRL Amateur Radio 9th Computer Networking Conference*, pp. 134–140, ARRL, 1990.

[2] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang, "MACAW: A Media Access Protocol for Wireless LANs," in *Proc. of ACM SIGCOMM '94*, 1994.

[3] IEEE Computer Society LAN MAN Standards Committee, ed., *IEEE Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*. IEEE Std 802.11-1997, The Institute of Electrical and Electronics Engineers, New York, 1997.

[4] J. J. Garcia-Luna-Aceves and C. L. Fullmer, "Floor Acquisition Multiple Access (FAMA) in Single-channel Wireless Networks," *ACM/Baltzer Mobile Networks and Applications*, vol. 4, no. 3, pp. 157–174, 1999.

[5] F. Talucci and M. Gerla, "MACA-BI (MACA by Invitation): A Receiver Oriented Access Protocol for Wireless Multihop Networks," in *Proc. of PIMRC '97*, 1997.

[6] J. J. Garcia-Luna-Aceves and A. Tzamaloukas, "Receiver-initiated Collision Avoidance in Wireless Networks," *ACM Wireless Networks*, vol. 8, pp. 249–263, 2002.

[7] F. A. Tobagi and L. Kleinrock, "Packet Switching in Radio Channels: Part II - the Hidden Terminal Problem in Carrier Sense Multiple-access Modes and the Busy-tone Solution," *IEEE Trans. on Communications*, vol. 23, no. 12, pp. 1417–1433, 1975.

[8] T. Ozugur, M. Naghshineh, P. Kermani, C. M. Olsen, B. Rezvani, and J. A. Copeland, "Balanced Media Access Methods for Wireless Networks," in *Proc. of ACM/IEEE MOBICOM '98*, pp. 21–32, Oct. 1998.

[9] T. Nandagopal, T. Kim, X. Gao, and V. Bharghavan, "Achieving MAC Layer Fairness in Wireless Packet Networks," in *ACM Mobicom 2000*, (Boston, MA, USA), Aug. 2000.

[10] N. H. Vaidya, P. Bahl, and S. Gupta, "Distributed Fair Scheduling in a Wireless LAN," in *ACM Mobicom 2000*, (Boston, MA, USA), Aug. 2000.

[11] H. Luo, S. Lu, and V. Bharghavan, “A New Model for Packet Scheduling in Multihop Wireless Networks,” in *ACM Mobicom 2000*, (Boston, MA, USA), Aug. 2000.

[12] Y. Wang and B. Bensaou, “Achieving Fairness in IEEE 802.11 DFWMAC with Variable Packet Length,” in *Proc. of IEEE GLOBECOM*, (San Antonio, Texas, U.S.A.), Nov. 2001.

[13] X. Huang and B. Bensaou, “On Max-min Fairness and Scheduling in Wireless Ad-Hoc Networks: Analytical Framework and Implementation,” in *ACM MobiHoc ’01*, 2001.

[14] H. Luo and S. Lu, “A Topology-Independent Fair Queueing Model in Ad Hoc Wireless Networks,” in *IEEE ICNP 2000*, (Osaka, Japan), Nov. 2000.

[15] H. Luo, P. Medvedev, J. Cheng, and S. Lu, “A Self-Coordinating Approach to Distributed Fair Queueing in Ad Hoc Wireless Networks,” in *IEEE INFOCOM 2001*, Apr. 2001.

[16] X. Zeng, R. Bagrodia, and M. Gerla, “GloMoSim: a Library for Parallel Simulation of Large-scale Wireless Networks,” in *Proc. of the 12th Workshop on Parallel and Distributed Simulations*, May 1998.

Authors' Biography

Yu Wang Yu Wang received the B.Eng. degree in Electronic Engineering from Shanghai Jiaotong University, Shanghai, China, in 1998, the M.Eng. degree in Electrical Engineering from National University of Singapore, Singapore, in 2001, and the M.Sc degree in Computer Engineering from University of California, Santa Cruz (UCSC), CA, U.S.A. in 2001. Since 2000, he has been a Ph.D. student in the Computer Communication Research Group (CCRG) at the Computer Engineering Department, UCSC. His research interests include Medium Access Control and Quality of Service in wireless ad-hoc networks. He is a student member of the IEEE.

J. J. Garcia-Luna-Aceves J. J. Garcia-Luna-Aceves received the B.S. degree in electrical engineering from the Universidad Iberoamericana, Mexico City, Mexico, in 1977, and the M.S. and Ph.D. degrees in electrical engineering from the University of Hawaii, Honolulu, HI, in 1980 and 1983, respectively. He is Professor of Computer Engineering at the University of California, Santa Cruz (UCSC). He directs the Computer Communication Research Group (CCRG), which is part of the Information Technologies Institute of the Baskin School of Engineering at UCSC.

He has been a Visiting Professor at Sun Laboratories and a consultant on protocol design for Nokia. Prior to joining UCSC in 1993, he was a Center Director at SRI International (SRI) in Menlo Park, California.

Dr. Garcia-Luna-Aceves has published a book and more than 225 refereed papers. He has been Program Co-Chair of ACM MobiHoc 2002 and ACM Mobicom 2000; Chair of the ACM SIG Multimedia; General Chair of ACM Multimedia '93 and ACM SIGCOMM '88; and Program Chair of IEEE MULTIMEDIA '92, ACM SIGCOMM '87, and ACM SIGCOMM '86. He has served in the IEEE Internet Technology Award Committee, the IEEE Richard W. Hamming Medal Committee, and the National Research Council Panel on Digitization and Communications Science of the Army Research Laboratory Technical Assessment Board. He has been on the editorial boards of the IEEE/ACM Transactions on Networking, the Multimedia Systems Journal, and the Journal of High Speed Networks. He received the SRI International Exceptional-Achievement Award in 1985 and 1989, and is a senior member of the IEEE.

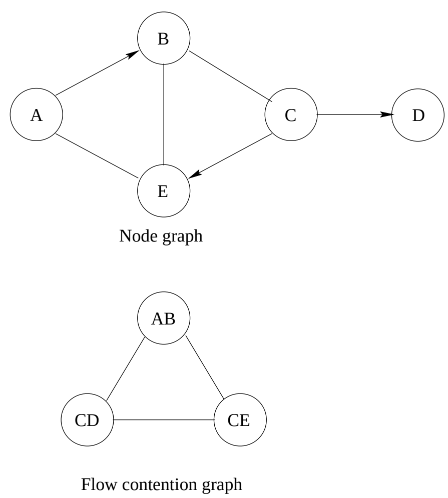

Figure 1: A simple network: node graph and flow contention graph

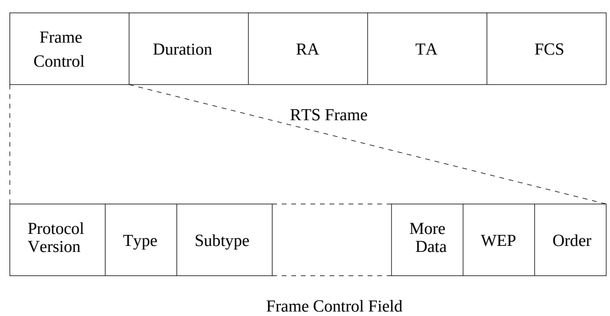

Figure 2: Illustration of the IEEE 802.11 frame structure

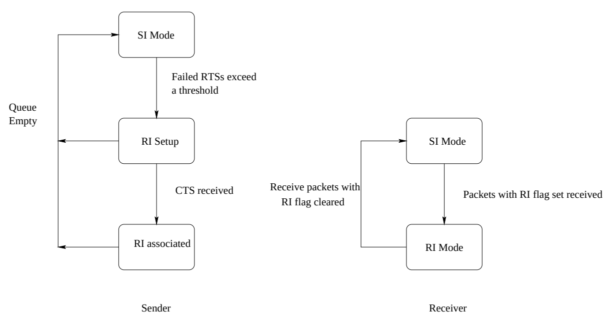

Figure 3: State transition diagram of sending and receiving nodes

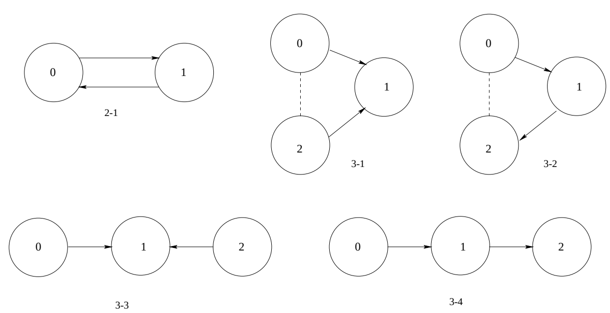

Figure 4: Networks with 2 or 3 nodes

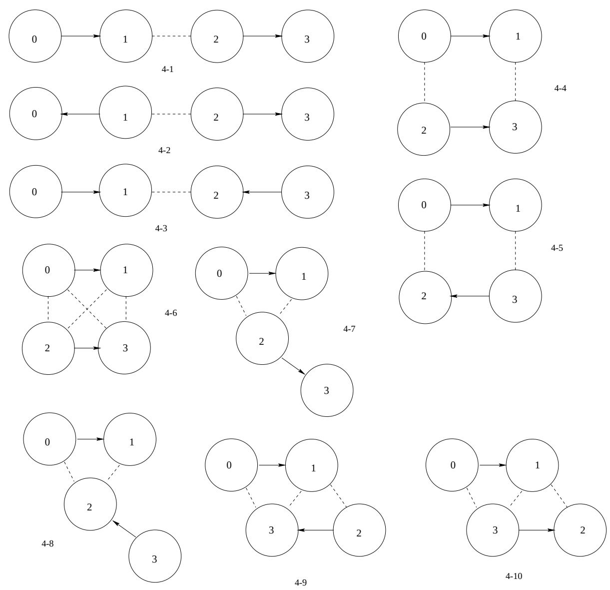

Figure 5: Networks with 4 nodes

Table 1: IEEE 802.11 protocol configuration parameters

RTS	CTS	data	ACK	DIFS	SIFS
20-byte	14-byte	1460-byte	14-byte	50μsec	10μsec

contention window	slot time	sync. time	prop. delay
31–1023	20μsec	192μsec	1μsec

Table 2: Fairness problems in the original IEEE 802.11 – two CBR flows

Config #	Flow #	Throughput (bps)	Flow #	Throughput (bps)	Aggregate (bps)
2-1	$0 \rightarrow 1$	8.06e+05	$1 \rightarrow 0$	7.99e+05	1.60e+06
3-1	$0 \rightarrow 1$	8.06e+05	$2 \rightarrow 1$	7.97e+05	1.60e+06
3-2	$0 \rightarrow 1$	7.97e+05	$1 \rightarrow 2$	8.07e+05	1.60e+06
3-3	$0 \rightarrow 1$	7.61e+05	$2 \rightarrow 1$	7.83e+05	1.54e+06
3-4	$0 \rightarrow 1$	7.69e+05	$1 \rightarrow 2$	8.39e+05	1.61e+06
4-1	$0 \rightarrow 1$	8.34e+04	$2 \rightarrow 3$	1.50e+06	1.58e+06
4-2	$1 \rightarrow 0$	8.20e+05	$2 \rightarrow 3$	8.14e+05	1.63e+06
4-3	$0 \rightarrow 1$	6.88e+05	$3 \rightarrow 2$	7.09e+05	1.40e+06
4-4	$0 \rightarrow 1$	8.24e+05	$2 \rightarrow 3$	8.08e+05	1.63e+06
4-5	$0 \rightarrow 1$	8.08e+05	$3 \rightarrow 2$	7.95e+05	1.60e+06
4-6	$0 \rightarrow 1$	8.07e+05	$2 \rightarrow 3$	7.95e+05	1.60e+06
4-7	$0 \rightarrow 1$	7.83e+05	$2 \rightarrow 3$	8.24e+05	1.61e+06
4-8	$0 \rightarrow 1$	1.55e+06	$3 \rightarrow 2$	2.81e+04	1.58e+06
4-9	$0 \rightarrow 1$	7.34e+05	$2 \rightarrow 3$	8.09e+05	1.54e+06
4-10	$0 \rightarrow 1$	7.81e+05	$3 \rightarrow 2$	8.26e+05	1.61e+06

Table 3: Throughput comparison for the IEEE 802.11, the measurement-based fair scheme (MFS) and the hybrid scheme (with RI mode) – two CBR flows

Config #	Scheme	Flow #	Throughput (bps)	Flow #	Throughput (bps)	Aggregate (bps)
3-3	802.11	$0 \rightarrow 1$	7.61e+05	$2 \rightarrow 1$	7.83e+05	1.54e+06
	+RImode	$0 \rightarrow 1$	7.94e+05	$2 \rightarrow 1$	7.74e+05	1.61e+06
	+MFS	$0 \rightarrow 1$	4.72e+05	$2 \rightarrow 1$	4.71e+05	9.43e+05
4-1	802.11	$0 \rightarrow 1$	8.34e+04	$2 \rightarrow 3$	1.50e+06	1.58e+06
	+RImode	$0 \rightarrow 1$	3.69e+05	$2 \rightarrow 3$	1.23e+06	1.60e+06
	+MFS	$0 \rightarrow 1$	9.79e+05	$2 \rightarrow 3$	5.34e+05	1.51e+06
4-3	802.11	$0 \rightarrow 1$	6.88e+05	$3 \rightarrow 2$	7.09e+05	1.40e+06
	+RImode	$0 \rightarrow 1$	6.65e+05	$3 \rightarrow 2$	6.43e+05	1.31e+06
	+MFS	$0 \rightarrow 1$	6.91e+05	$3 \rightarrow 2$	6.98e+05	1.39e+06
4-8	802.11	$0 \rightarrow 1$	1.55e+06	$3 \rightarrow 2$	2.81e+04	1.58e+06
	+RImode	$0 \rightarrow 1$	1.28e+06	$3 \rightarrow 2$	3.19e+05	1.60e+06
	+MFS	$0 \rightarrow 1$	5.22e+05	$3 \rightarrow 2$	9.86e+05	1.51e+06
4-9	802.11	$0 \rightarrow 1$	7.34e+05	$2 \rightarrow 3$	8.09e+05	1.54e+06
	+RImode	$0 \rightarrow 1$	8.15e+05	$2 \rightarrow 3$	7.42e+05	1.56e+06
	+MFS	$0 \rightarrow 1$	4.72e+05	$2 \rightarrow 3$	4.71e+05	9.42e+05

Table 4: Fairness problems in the original IEEE 802.11 – two FTP flows

Config #	Flow #	Throughput (bps)	Flow #	Throughput (bps)	Aggregate (bps)
2-1	$0 \rightarrow 1$	4.66e+05	$1 \rightarrow 0$	4.68e+05	9.34e+05
3-1	$0 \rightarrow 1$	4.72e+05	$2 \rightarrow 1$	4.73e+05	9.45e+05
3-2	$0 \rightarrow 1$	4.56e+05	$1 \rightarrow 2$	4.79e+05	9.35e+05
3-3	$0 \rightarrow 1$	4.92e+05	$2 \rightarrow 1$	3.84e+05	8.75e+05
3-4	$0 \rightarrow 1$	3.52e+05	$1 \rightarrow 2$	5.48e+05	9.00e+05
4-1	$0 \rightarrow 1$	0	$2 \rightarrow 3$	9.26e+05	9.29e+05
4-2	$1 \rightarrow 0$	(4.88±1.03)e+05	$2 \rightarrow 3$	(4.53±1.02)e+05	9.42e+05
4-3	$0 \rightarrow 1$	(5.30±4.32)e+05	$3 \rightarrow 2$	(3.92±4.38)e+05	9.22e+05
4-4	$0 \rightarrow 1$	4.49e+05	$2 \rightarrow 3$	4.36e+05	8.84e+05
4-5	$0 \rightarrow 1$	4.75e+05	$3 \rightarrow 2$	4.74e+05	9.49e+05
4-6	$0 \rightarrow 1$	4.75e+05	$2 \rightarrow 3$	4.74e+05	9.49e+05
4-7	$0 \rightarrow 1$	9.28e+05	$2 \rightarrow 3$	0	9.30e+05
4-8	$0 \rightarrow 1$	9.29e+05	$3 \rightarrow 2$	0	9.30e+05
4-9	$0 \rightarrow 1$	4.27e+05	$2 \rightarrow 3$	4.49e+05	8.76e+05
4-10	$0 \rightarrow 1$	3.76e+05	$3 \rightarrow 2$	5.26e+05	9.02e+05

Table 5: Throughput comparison for the IEEE 802.11, the measurement-based fair scheme (MFS) and the hybrid scheme (with RI mode) – two FTP flows

Config #	Scheme	Flow #	Throughput (bps)	Flow #	Throughput (bps)	Aggregate (bps)
3-4	802.11	$0 \rightarrow 1$	3.52e+05	$1 \rightarrow 2$	5.48e+05	9.00e+05
	+RImode	$0 \rightarrow 1$	3.30e+05	$1 \rightarrow 2$	5.54e+05	8.85e+05
	+MFS	$0 \rightarrow 1$	1.49e+05	$1 \rightarrow 2$	1.76e+05	3.25e+05
4-1	802.11	$0 \rightarrow 1$	0	$2 \rightarrow 3$	9.26e+05	9.29e+05
	+RImode	$0 \rightarrow 1$	-	$2 \rightarrow 3$	-	-
	+MFS	$0 \rightarrow 1$	2.16e+05	$2 \rightarrow 3$	2.28e+05	4.44e+05
4-2	802.11	$1 \rightarrow 0$	(4.88 ± 1.03)e+05	$2 \rightarrow 3$	(4.53 ± 1.02)e+05	9.42e+05
	+RImode	$0 \rightarrow 1$	(4.39 ± 0.99)e+05	$3 \rightarrow 2$	(5.02 ± 0.98)e+05	9.40e+05
	+MFS	$0 \rightarrow 1$	2.49e+05	$3 \rightarrow 2$	2.52e+05	5.02e+05
4-3	802.11	$0 \rightarrow 1$	(5.30 ± 4.32)e+05	$3 \rightarrow 2$	(3.92 ± 4.38)e+05	9.22e+05
	+RImode	$0 \rightarrow 1$	(3.97 ± 0.71)e+05	$3 \rightarrow 2$	(4.55 ± 0.78)e+05	8.52e+05
	+MFS	$0 \rightarrow 1$	2.20e+05	$3 \rightarrow 2$	2.20e+05	4.41e+05
4-7	802.11	$0 \rightarrow 1$	9.28e+05	$2 \rightarrow 3$	0	9.30e+05
	+RImode	$0 \rightarrow 1$	-	$2 \rightarrow 3$	-	-
	+MFS	$0 \rightarrow 1$	4.24e+05	$2 \rightarrow 3$	4.04e+05	8.28e+05
4-8	802.11	$0 \rightarrow 1$	9.29e+05	$3 \rightarrow 2$	0	9.30e+05
	+RImode	$0 \rightarrow 1$	-	$3 \rightarrow 2$	-	-
	+MFS	$0 \rightarrow 1$	2.46e+05	$3 \rightarrow 2$	2.06e+05	4.52e+05